CLAIMS

1. A method for static rate flow control, comprising:

receiving a sequence of data packets for transmission over a network, including at least first and second packets having a common destination address on the network, the first and second packets having respective first and second lengths;

transmitting the first packet to the destination address;

responsive to transmitting the first packet, placing an entry in a flow control table;

setting a timeout period for the entry responsive to the first length; and

transmitting the second packet only after the timeout period has expired.

2. A method according to claim 1, wherein setting the timeout period comprises setting the period to a multiple of a time interval required to transmit the first packet, which time interval is determined responsive to the first length.

3. A method according to claim 2, wherein the multiple is a number greater than one.

4. A method according to claim 3, wherein transmitting the first and second packets comprises transmitting the packets through a network port having a port transmission rate, while the network is configured to pass the packets to the destination address at a destination reception rate that is less than the port transmission rate, and wherein the multiple is determined responsive to a relation of the port transmission rate and the destination reception rate.

5. A method according to claim 4, wherein the multiple is substantially equal to a ratio of the port transmission rate to the destination reception rate.

6. A method according to claim 1, wherein placing the entry in the flow control table comprises placing the entry corresponding to the destination address of the first packet, and comprising removing the entry from the table upon expiration of the timeout period.

7. A method according to claim 6, wherein transmitting the second packet comprises checking the table with reference to the destination address of the second packet to determine whether the table contains the entry, and sending the second packet only when the entry is absent from the table.

8. A method according to claim 7, wherein transmitting the second packet comprises restoring the entry corresponding to the destination address to the table upon transmitting the second packet.

9. A method according to claim 1, wherein placing the entry comprises placing the entry corresponding to the destination address of the first packet among a plurality of entries in the table corresponding to different destination addresses in the network to which the packets in the sequence are directed.

10. A method according to claim 9, wherein some of the destination addresses in the network are subject to static flow control, while others are not, and wherein placing the entry comprises associating an attribute with the first packet indicating whether the destination address of the first packet is subject to the static flow

control, and placing the entry in the table responsive to the attribute.

11. A method according to claim 9, wherein placing the entry comprises writing the entry in a memory having a size determined according to a maximum number of the different destination addresses to which static flow control scheduling is likely to be applied simultaneously.

12. A method according to claim 1, wherein the network comprises a switch fabric, and wherein the destination address comprises a Destination Local Identifier (DLID).

13. A network end-node device, for transmitting a sequence of data packets over a network, including at least first and second packets having a common destination address on the network, the first and second packets having respective first and second lengths, the device comprising:

a memory, configured to hold a flow control table; and

link output circuitry adapted to transmit the first packet to the destination address and, responsive to transmitting the first packet, to place an entry in the flow control table and to set a timeout period for the entry responsive to the first length, and to transmit the second packet only after the timeout period has expired.

14. A device according to claim 13, wherein the timeout period is set to a multiple of a time interval required to transmit the first packet, which time interval is determined responsive to the first length.

15. A device according to claim 14, wherein the multiple is a number greater than one.

16. A device according to claim 15, wherein the link output circuitry comprises a network port having a port transmission rate, while the network is configured to pass the packets to the destination address at a destination reception rate that is less than the port transmission rate, and wherein the multiple is determined responsive to a relation of the port transmission rate and the destination reception rate.

17. A device according to claim 16, wherein the multiple is substantially equal to a ratio of the port transmission rate to the destination reception rate.

18. A device according to claim 13, wherein the entry in the flow control table corresponds to the destination address of the first packet, and wherein the link output circuitry is adapted to remove the entry from the table upon expiration of the timeout period.

19. A device according to claim 18, wherein the link output circuitry is adapted to check the table with reference to the destination address of the second packet to determine whether the table contains the entry, and to send the second packet only when the entry is absent from the table.

20. A device according to claim 19, wherein the link output circuitry is adapted to restore the entry corresponding to the destination address to the table upon transmitting the second packet.

21. A device according to claim 13, wherein the link output circuitry is adapted to place a plurality of

entries in the table corresponding to different destination addresses in the network to which the packets in the sequence are directed, including the entry in the flow control table corresponding to the destination address of the first packet.

22. A device according to claim 21, wherein some of the destination addresses in the network are subject to static flow control, while others are not, and wherein an attribute is associated with each of the packets indicating whether the destination addresses of the packets are subject to the static flow control, and wherein the link output circuitry is adapted to place each of the entries in the table responsive to the attribute.

23. A device according to claim 21, wherein the memory has a size allocated to the table that is determined according to a maximum number of the different destination addresses to which static flow control scheduling is likely to be applied simultaneously.

24. A device according to claim 13, wherein the network comprises a switch fabric, and the device comprises a channel adapter, and wherein the destination address comprises a Destination Local Identifier (DLID).